arXiv:1701.00056v1 [math.ST] 31 Dec 2016

Compressed sensing and optimal denoising of monotone signals

Eftychios A. Pnevmatikakis

Center for Computational Biology, Flatiron Institute, Simons Foundation, New York, NY 10010

Abstract

We consider the problems of compressed sensing and optimal denoising for signals $\boldsymbol{x_0} \in \mathbb{R}^N$ that are monotone, i.e., $\boldsymbol{x_0}(i+1) \geq \boldsymbol{x_0}(i)$, and sparsely varying, i.e., $\boldsymbol{x_0}(i+1) > \boldsymbol{x_0}(i)$ only for a small number k of indices i. We approach the compressed sensing problem by minimizing the total variation norm restricted to the class of monotone signals subject to equality constraints obtained from a number of measurements $A\boldsymbol{x_0}$. For random Gaussian sensing matrices $A \in \mathbb{R}^{m \times N}$ we derive a closed form expression for the number of measurements m required for successful reconstruction with high probability. We show that the probability undergoes a phase transition as m varies, and depends not only on the number of change points, but also on their location. For denoising we regularize with the same norm and derive a formula for the optimal regularizer weight that depends only mildly on $\boldsymbol{x_0}$. We obtain our results using the statistical dimension tool.

1 Introduction

We consider N-dimensional signals $\boldsymbol{x_0}$ that are sparse in a certain basis. We are interested in the following problems

$$\min_{\boldsymbol{x}} f(\boldsymbol{x}), \text{ subject to } A\boldsymbol{x} = A\boldsymbol{x_0}, \tag{CS}$$

and

$$\min_{\boldsymbol{x}} \frac{1}{2}\|\boldsymbol{y} - \boldsymbol{x}\|^2 + \lambda f(\boldsymbol{x}), \tag{DN}$$

with $\boldsymbol{y} = \boldsymbol{x_0} + \boldsymbol{\varepsilon}$ and $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma^2 I_N)$. Here $f : \mathbb{R}^N \mapsto \mathbb{R} \cup \{\infty\}$ is a convex function that characterizes the structure of $\boldsymbol{x_0}$. For the compressed sensing problem (CS) we are interested in deriving the minimum number of measurements m such that the solution of (CS) coincides with $\boldsymbol{x_0}$ with high probability for standard normal i.i.d. sensing matrices $A \in \mathbb{R}^{m \times N}$. For the denoising problem (DN) we are interested in calculating the minimax risk, i.e., the optimal value of (DN) minimized over λ for the worst case of noise power σ^2. The two quantities are closely related due to some recent results reviewed briefly next.

2 Basic tools

Definition 2.1 (Descent cones)**.** *The descent cone of a convex function $f : \mathbb{R}^N \mapsto \mathbb{R}$ at a point $\boldsymbol{x} \in \mathbb{R}^N$ is defined as the set of all non-increasing directions, i.e.,*

$$\mathcal{D}(f, \boldsymbol{x}) = \bigcup_{\tau>0} \{\boldsymbol{y} \in \mathbb{R}^N : f(\boldsymbol{x} + \tau \boldsymbol{y}) \leq f(\boldsymbol{x})\}.$$

Definition 2.2 (Statistical dimension (Amelunxen et al., 2014))**.** *The statistical dimension (SD) of a convex closed cone $\mathcal{C} \in \mathbb{R}^N$ is defined as*

$$\delta(\mathcal{C}) = \mathbb{E}_{\boldsymbol{g} \sim \mathcal{N}(\mathbf{0}, I_N)} \|\Pi_{\mathcal{C}}(\boldsymbol{g})\|^2,$$

where $\boldsymbol{g}$ is a standard Gaussian vector, and $\Pi_{\mathcal{C}}$ is the projection onto $\mathcal{C}$.

In a groundbreaking work, Amelunxen et al. (2014) shows that the SD of the descent cone at the true point $\boldsymbol{x_0}$, coincides with the phase transition curve (PTC) of the CS problem.

Theorem 2.3 (Phase transitions (Amelunxen et al., 2014))**.** *For an i.i.d. standard random Gaussian matrix $A \in \mathbb{R}^{m \times N}$ the convex problem* (CS) *succeeds with probability at least $1 - \exp(-t^2/4)$ if*

$$m \geq \delta(\mathcal{D}(f, \boldsymbol{x_0})) + t\sqrt{N},$$

and fails with probability at least $1 - \exp(-t^2/4)$ if

$$m \leq \delta(\mathcal{D}(f, \boldsymbol{x_0})) - t\sqrt{N}.$$

Furthermore, Amelunxen et al. (2014) shows that the SD can also be expressed as the expected distance from the subdifferential of f at $\boldsymbol{x_0}$:

$$\delta(\mathcal{D}(f, \boldsymbol{x_0})) = \mathbb{E}_{\boldsymbol{g} \sim \mathcal{N}(\mathbf{0}, I_N)} [\min_{\tau \geq 0} \operatorname{dist}(\boldsymbol{g}, \tau \partial f(\boldsymbol{x_0}))^2] \tag{1}$$

Theorem 2.4 (Minimax risk (Oymak and Hassibi, 2012))**.** *Let $\boldsymbol{x}^*(\lambda)$ the solution of the denoising problem* (DN) *with regularizer weight λ and let*

$$\eta_f(\boldsymbol{x_0}) = \min_{\lambda \geq 0} \max_{\sigma > 0} \frac{\mathbb{E}\|\boldsymbol{x}^*(\lambda) - \boldsymbol{x_0}\|^2}{\sigma^2},$$

the minimax risk for $\boldsymbol{x_0}$ over all possible σ. Then:

$$\eta_f(\boldsymbol{x_0}) = \min_{\tau \geq 0} \mathbb{E}_{\boldsymbol{g} \sim \mathcal{N}(\mathbf{0}, I)} [\operatorname{dist}(\boldsymbol{g}, \tau \partial f(\boldsymbol{x_0}))^2], \tag{2}$$

where $\boldsymbol{g}$ is a standard normal vector. Moreover the risk is maximized for $\sigma \to 0$ and if τ^* is the value that minimizes (2)*, then $\lambda^* = \tau^* \sigma$ is the optimal choice as $\sigma \to 0$.*

The similarity between (1) and (2) is striking and actually Amelunxen et al. (2014) proves that the two quantities are indeed close:

$$\delta(\mathcal{D}(f, \boldsymbol{x_0})) \leq \eta_f(\boldsymbol{x_0}) \leq \delta(\mathcal{D}(f, \boldsymbol{x_0})) + 2 \frac{\sup_{\boldsymbol{w} \in \partial f(\boldsymbol{x_0})} \|\boldsymbol{w}\|}{f(\boldsymbol{x_0}/\|\boldsymbol{x_0}\|)}.$$

3 Phase transitions for the recovery of sparsely varying monotone signals

We consider signals $\boldsymbol{x} \in \mathbb{R}^N$ that are increasing, i.e., $\boldsymbol{x}(i+1) \geq \boldsymbol{x}(i)$ and are sparsely varying, i.e, $\boldsymbol{x}(i+1) > \boldsymbol{x}(i)$ for a number of k indexes. A convex function that promotes this structure can be derived by restricting the total variation (TV) norm to the space of monotone signals:

$$f(\boldsymbol{x}) = \begin{cases} \boldsymbol{x}(N) - \boldsymbol{x}(1), & \boldsymbol{x}(i+1) \geq \boldsymbol{x}(i), i \in [N-1] \\ \infty, & \text{otherwise.} \end{cases} \tag{3}$$

where $[N] = \{1, 2, \ldots, N\}$. Our results rely heavily on the following calculation of the SDs of the cones induced by monotone signals, proven in Amelunxen et al. (2014, App. C.4).

Fact 3.1. *Let the cones*

$$\mathcal{C}_1^N = \{\boldsymbol{x} \in \mathbb{R}^N : \boldsymbol{x}(1) \leq \boldsymbol{x}(2) \leq \ldots \leq \boldsymbol{x}(N)\}$$
$$\mathcal{C}_2^N = \{\boldsymbol{x} \in \mathbb{R}^N : 0 \leq \boldsymbol{x}(1) \leq \boldsymbol{x}(2) \leq \ldots \leq \boldsymbol{x}(N)\}.$$

Then we have $\delta(\mathcal{C}_1^N) = H_N$, *and* $\delta(\mathcal{C}_2^N) = \frac{1}{2}H_N$, *where* $H_N = \sum_{i=1}^N \frac{1}{i}$, *denotes the* N*-th harmonic number.*

3.1 Computation of the statistical dimension

According to Theorem 2.3 to compute the PTC for the CS problem, we need to characterize $\mathcal{D}(f, \boldsymbol{x_0})$.

Lemma 3.2. *Let* $\Omega = \{i \in \{2, \ldots, N\} : \boldsymbol{x_0}(i) > \boldsymbol{x_0}(i-1)\}$ *and define* $i_1 < i_2 < \ldots < i_k$ *the elements of* Ω *in increasing order. The descent cone of the norm* f *of* (3) *at* $\boldsymbol{x_0}$ *is given by*

$$\mathcal{D}(f, \boldsymbol{x_0}) = \left\{\boldsymbol{y} \in \mathbb{R}^N : \begin{array}{c} \boldsymbol{y}(i_1) \leq \boldsymbol{y}(i_1+1) \leq \ldots \leq \boldsymbol{y}(i_2-1) \\ \vdots \\ \boldsymbol{y}(i_{k-1}) \leq \boldsymbol{y}(i_{k-1}+1) \leq \ldots \leq \boldsymbol{y}(i_k-1) \\ \boldsymbol{y}(i_k) \leq \ldots \leq \boldsymbol{y}(N) \leq \boldsymbol{y}(1) \leq \ldots \leq \boldsymbol{y}(i_1-1) \end{array} \right\}. \tag{4}$$

Proof. From Definition 2.1, $\boldsymbol{y} \in \mathcal{D}(f, \boldsymbol{x_0})$ if there exists $\tau > 0$, such that $\boldsymbol{x_0} + \tau\boldsymbol{y}$ is monotone, and $f(\boldsymbol{x_0} + \tau\boldsymbol{y}) \leq f(\boldsymbol{x_0})$:

$$\begin{aligned} \boldsymbol{x_0}(N) + \tau\boldsymbol{y}(N) - \boldsymbol{x_0}(1) - \tau\boldsymbol{y}(1) \leq \boldsymbol{x_0}(N) - \boldsymbol{x_0}(1) \Rightarrow \\ \boldsymbol{y}(N) \leq \boldsymbol{y}(1). \end{aligned}$$

For the monotonicity of $\boldsymbol{x_0} + \tau\boldsymbol{y}$, we consider two cases: If $i \notin \Omega$, then $\boldsymbol{x_0}(i) = \boldsymbol{x_0}(i-1)$, and

$$\boldsymbol{x_0}(i) + \tau\boldsymbol{y}(i) \geq \boldsymbol{x_0}(i-1) + \tau\boldsymbol{y}(i-1) \Rightarrow \boldsymbol{y}(i) \geq \boldsymbol{y}(i-1).$$

If $i \in \Omega$, then $\boldsymbol{x_0}(i) > \boldsymbol{x_0}(i-1)$ and $\boldsymbol{y}(i)$ can be chosen arbitrarily since there is always a small enough τ that will preserve monotonicity. Combining everything we get (4). □

Lemma 3.2 states that the descent cone $\mathcal{D}(f, \boldsymbol{x_0})$ can be expressed as the product of k disjoint convex cones of monotonically increasing signals. Using Fact 3.1, we derive the following simple formula for $\delta(\mathcal{D}(f, \boldsymbol{x_0}))$ as the sum of the SDs of the simpler disjoint cones.

Theorem 3.3. *Let $\Omega = \{i \in \{2, \dots, N\} : \boldsymbol{x_0}(i) > \boldsymbol{x_0}(i-1)\}$ and define $i_1 < i_2 < \ldots < i_k$ the elements of Ω in increasing order. The SD of the descent cone at $\boldsymbol{x_0}$ equals*

$$\delta(\mathcal{D}(f, \boldsymbol{x_0})) = \sum_{j=2}^{k} H_{i_j - i_{j-1}} + H_{N+i_1-i_k}. \tag{5}$$

3.2 Dependence on the change points location

The closed form of the SD allows for a characterization of the worst case analysis for a given number of variations k. These locations have to occur periodically every N/k steps, with $i_1 \approx N/2k$. If $r_{N,k} = \mathrm{mod}(N, k)$, then the SD becomes $(k - r_{N,k})H_{[N/k]} + r_{N,k}H_{[N/k]+1}.$, where $[\cdot]$ here denotes the integer part. For moderately large N/k this converges to $kH_{N/k} \to k(\log(N/k) + \gamma)$, where $\gamma \approx 0.577$ is the Euler-Mascheroni constant.

Similarly, the best case occurs when all change points occur consecutively. In this case the SD becomes $(k-1) + H_{N+1-k}$. What is perhaps of most interest is the average SD under certain distribution assumptions of the k change points. We can asymptotically compute this in the case where these k points are distributed uniformly at random.

Theorem 3.4. *Assume that the k change points are chosen uniformly at random and let $N, k \to \infty$ with $k/N = \varepsilon$, $0 < \varepsilon < 1$. Define $\delta_U(\varepsilon)$ the normalized (divided by the ambient dimension) SD averaged over all possible choices of $k = \varepsilon N$ "jump" points. Then we have*

$$\delta_U(\varepsilon) = \frac{\varepsilon \log(1/\varepsilon)}{1-\varepsilon}. \tag{6}$$

Proof. Let $i_1 < i_2 < \ldots < i_k$ the change points selected uniformly randomly and define the sequence of lengths $l_j = i_{j+1} - i_j$ for $j \in [k-1]$ and $l_k = N - i_k + i_1$. When $N, k \to \infty$ the distribution of each l_j converges to a geometric distribution with parameter $\varepsilon = k/N$. Then we have

$$\begin{aligned}
\delta_U(\varepsilon) &= \lim_{N\to\infty} \frac{1}{N} \mathbb{E}\left[\sum_{j=1}^{k} H_{l_j}\right] = \varepsilon \mathbb{E}\left[H_{l_j}\right] \\
&= \varepsilon^2 \sum_{n=1}^{\infty} H_n (1-\varepsilon)^{n-1} = \frac{\varepsilon^2}{1-\varepsilon} \sum_{n=1}^{\infty} \frac{1}{n} \sum_{m=n}^{\infty} (1-\varepsilon)^m \\
&= \frac{\varepsilon^2}{1-\varepsilon} \sum_{n=1}^{\infty} \frac{1}{n} \frac{(1-\varepsilon)^n}{\varepsilon} = \frac{\varepsilon \log(1/\varepsilon)}{1-\varepsilon}. \quad \square
\end{aligned}$$

Fig. 1 shows the three different cases for the SD, and illustrates its dependence on the location of the jump points.

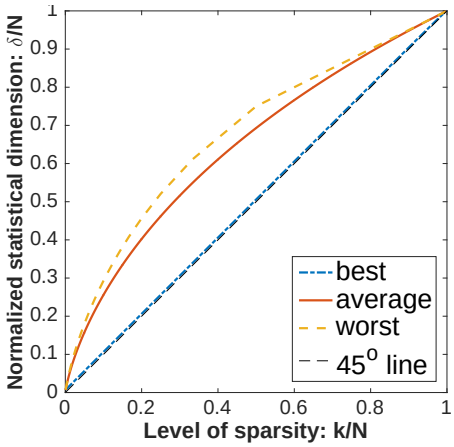

Figure 1: Behavior of the SD as a function of the degree sparsity and location of change points. The best (blue, dash-dot), average (red, solid), and worst (yellow, dashed) cases are shown.

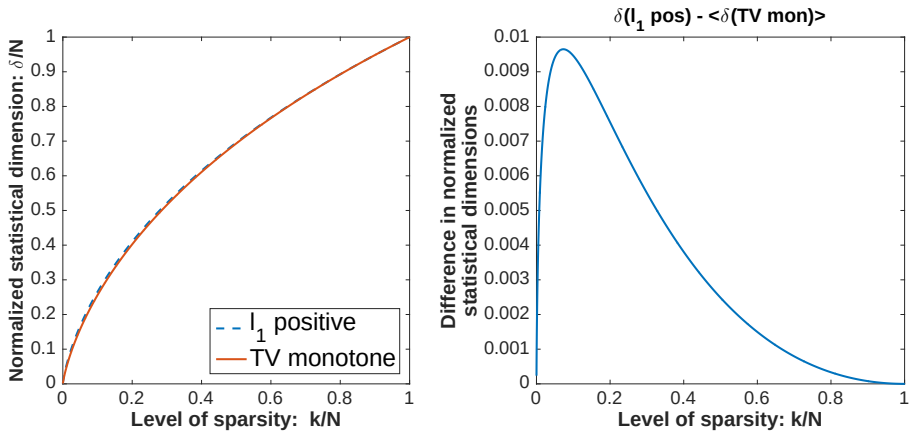

Figure 2: Comparison of the SD for monotone sparsely varying signals and sparse non-negative signals. The SD for reconstructing sparse non-negative signals (red, solid) as computed in Donoho et al. (2009) compared to the asymptotic average limit of (6) (blue, dashed). Somewhat surprisingly, the two curves almost coincide with the positive l_1 norm PTC being always larger than the PTC of for the monotone signals for the same number of variations. The difference is plot in the right panel.

In Fig. 2 we plot the SD as computed in (6) (blue), compared to the PTC for the reconstruction of sparse non-negative signals (red) as this is computed in Donoho et al. (2009), using the l_1 norm restricted to non-negative signals as the structure inducing function f and solving (CS). The two curves are very near, although the PTC curve for sparse non-negative signals is slightly larger. The difference between the two different curves (Fig. 2 right) attains a maximum of ≈ 0.0096 for $k/N \approx 0.0731$. We examined this difference in practice: For the sparse non-negative signal we considered signals $\boldsymbol{x_0} \in \mathbb{R}^N, N = 1000$, with $k = 73$ non-zero entries. Then random Gaussian matrices $A \in \mathbb{R}^{m\times N}$ were constructed, with $m = 201, \ldots, 220$, and we tried to reconstruct $\boldsymbol{x_0}$ from the samples $A\boldsymbol{x_0}$ by solving (CS). We solved the same problem also for the case of sparsely varying increasing signals, where now $k = 73$, refers to the number of change points, chosen uniformly at random. For each m we performed 300 iterations, and the reconstruction $\hat{\boldsymbol{x}}$ was deemed successful if $\|\hat{\boldsymbol{x}} - \boldsymbol{x_0}\|/\|\boldsymbol{x_0}\| \leq 10^{-4}$. The results show that the probability of accurate reconstruction crosses 50% within one measurement from the point predicted by the theoretical calculation of the SD, and that the difference of measurements required for 50% reconstruction probability is around 10 measurements, as predicted by the difference of the two SDs (data not shown due to space constraints). These simulation results validate the theoretical analysis.

3.3 The case of non-negative monotone signals

We also consider the case of increasing and sparsely varying signals $\boldsymbol{x} \in \mathbb{R}^N$, that are also non-negative, which we denote without loss of generality as $\boldsymbol{x}(0) = 0$, and consider the first entry as a change point if $\boldsymbol{x}(1) > 0$. In this case we consider the following convex regularizer $f(\boldsymbol{x}) = \boldsymbol{x}(N)$ for monotonically increasing signals and $f(\boldsymbol{x}) = \infty$ otherwise. Using a similar procedure we can derive the descent cone $\mathcal{D}(f, \boldsymbol{x_0})$ and from Fact 3.1 get a similar formula for the SD:

Theorem 3.5. *Let $\Omega = \{i \in [N] : \boldsymbol{x_0}(i) > \boldsymbol{x_0}(i-1)\}$ and define $i_1 < i_2 < \ldots < i_k$ the elements of Ω in increasing order. Then the SD of the descent cone at $\boldsymbol{x_0}$ is given by*

$$\delta(\mathcal{D}(f, \boldsymbol{x_0})) = \frac{1}{2}H_{i_1-1} + \frac{1}{2}H_{N+1-i_k} + \sum_{j=2}^{k} H_{i_j - i_{j-1}}.$$

4 Optimal denoising

For the case of monotone, sparsely varying, non-negative signals it is also possible to compute the minimax denoising risk, by using Theorem 2.4. To consider the risk of the denoising problem (DN), we first derive the subdifferential of f. Let G be the $N \times N$ matrix with $[G]_{ij} = 1_{\{i=j\}} - 1_{\{i=j+1\}}$ and define the function $h : \mathbb{R}^N \mapsto \mathbb{R} \cup \{\infty\}$, with

$$h(\boldsymbol{z}) = \begin{cases} \mathbf{1}^\top \boldsymbol{z}, & \boldsymbol{z}(i) \geq 0, i \in [N] \\ \infty, & \text{otherwise.} \end{cases}$$

Then $f(\boldsymbol{x}) = h(G\boldsymbol{x})$ and $\partial f(\boldsymbol{x}) = G^\top \partial h(G\boldsymbol{x})$ and

$$\partial f(\boldsymbol{x}) = G^\top \boldsymbol{w}, \text{ with } \begin{cases} \boldsymbol{w}(i) = 1, & \boldsymbol{x}(i) > \boldsymbol{x}(i-1) \\ \boldsymbol{w}(i) \leq 1, & \boldsymbol{x}(i) = \boldsymbol{x}(i-1) \end{cases}.$$

Therefore the distance of any vector $\boldsymbol{g} \in \mathbb{R}^N$ from $\partial f(\boldsymbol{x_0})$ can be computed by solving the following quadratic program

$$\begin{aligned} &\underset{\boldsymbol{w},\tau}{\text{minimize}} \ \|\boldsymbol{g} - G^\top \boldsymbol{w}\|^2, \\ &\text{subject to: } \tau \geq 0, \{w(j) = \tau, j \in \Omega\}, \{w(j) \leq \tau, j \in \Omega^c\}. \end{aligned} \tag{QP}$$

Lemma 4.1. *Consider the quadratic program* (QP) *and let i_k denote the last element of Ω. Then the optimal τ^* is given by*

$$\tau^* = \max\left(\max_{j=i_k,\ldots,N} \left\{ \sum_{n=j}^{N} \boldsymbol{g}(n) \right\}, 0 \right). \tag{7}$$

Proof. We consider the Lagrangian function

$$\mathcal{L}(\boldsymbol{w}, \tau, \boldsymbol{\lambda}, \lambda_\tau) = \frac{1}{2}\|\boldsymbol{g} - G^\top \boldsymbol{w}\|^2 - \lambda_\tau \tau + \boldsymbol{\lambda}^\top (\boldsymbol{w} - \tau \mathbf{1}_N^\top). \tag{8}$$

The dual variable constraints and the first order optimality conditions of (QP) can be written as

$$\begin{aligned} GG^\top \boldsymbol{w} - G\boldsymbol{g} + \boldsymbol{\lambda} &= \mathbf{0}, && (9)\\ \mathbf{1}^\top \boldsymbol{\lambda} + \lambda_\tau &= 0, && (10)\\ \boldsymbol{\lambda}(j) \geq 0, \quad \boldsymbol{w}(j) &\leq \tau, \quad j \notin \Omega && (11)\\ \boldsymbol{\lambda}(j)(\boldsymbol{w}(j) - \tau) &= 0, \quad j \notin \Omega && (12)\\ \boldsymbol{w}(j) &= \tau, \quad j \in \Omega && (13)\\ \lambda_\tau \geq 0, \quad \lambda_\tau \tau &= 0. && (14) \end{aligned}$$

From (9)
$$\boldsymbol{w} = \underbrace{(GG^\top)^{-1}G}_{E}\boldsymbol{g} - \underbrace{(GG^\top)^{-1}}_{F}\boldsymbol{\lambda}, \tag{15}$$

where the matrices E, F can be computed explicitly: $[E]_{ij} = 1_{\{j \geq i\}}$ and $[F]_{ij} = N - \max\{i,j\} + 1$, and using (10) gives

$$\boldsymbol{w}(j) = \sum_{i=j}^{N} \boldsymbol{g}(i) + (N - j + 1)\lambda_\tau + \sum_{i=j+1}^{N} (i-j)\boldsymbol{\lambda}(i), \tag{16}$$

for $j \in [N]$. Now suppose let i_k the last change point and suppose that $M = \max_{j=i_k,\ldots,N} \left\{ \sum_{n=j}^{N} \boldsymbol{g}(n) \right\}$. Consider first the case where $M < 0$ and suppose that $\tau > 0$. In this case from (14) we have $\lambda_\tau = 0$. Plugging this into (16) for $j = N$ we get $\boldsymbol{w}(N) = \boldsymbol{g}(N) < 0 \Rightarrow \boldsymbol{w}(N) - \tau < 0 \overset{(12)}{\Rightarrow} \boldsymbol{\lambda}(N) = 0$. Decreasing j and proceeding similarly we get $\boldsymbol{\lambda}(N) = \boldsymbol{\lambda}(N-1) = \ldots = \boldsymbol{\lambda}(i_k + 1) = 0$. Now for $j = i_k$ we get $\boldsymbol{w}(i_k) = \sum_{j=i_k}^{N} \boldsymbol{g}(j) < 0$, and (13) cannot be satisfied for $\tau > 0$. Therefore $\tau = 0$.

Now assume that $M > 0$, and that this maximum occurs at the location $N - l$. Then by plugging $j = N - l$ into (16) and the nonnegativity of the dual variables we have that $\boldsymbol{w}(N - l) \geq$

$M \Rightarrow \tau \geq M \Rightarrow \lambda_\tau = 0$. We proceed as before: For $j = N$ (16) gives $\boldsymbol{w}(N) < M \leq \tau \Rightarrow \boldsymbol{\lambda}(N) = 0$. And similarly $\boldsymbol{\lambda}(N) = \boldsymbol{\lambda}(N-1) = \ldots = \boldsymbol{\lambda}(N-l+1) = 0$. Plugging this into (16) for $j = N - l$ we get $\boldsymbol{w}(N-l) = M$. Since $\boldsymbol{w}(N-l) \leq \tau$ we get that $\tau = M$. □

Lemma 4.1 allows us to estimate the regularizer that minimizes (2) by estimating τ^{avg} that arises in (1), and consequently set the regularizer $\lambda = \tau^{avg}\sigma$. In general $\tau^{avg} = M(N - i_k + 1)$ where $M(n)$ is the expected value of the maximum of a standard Gaussian random walk of n steps, truncated at 0. $M(1) = 1/\sqrt{2\pi}$, and $M(2) = (1+\sqrt{2})/(2\sqrt{\pi})$. In general $M(n)$ cannot be computed explicitly, but can be easily upper bounded: Let $X_i \sim \mathcal{N}(0,1)$, $S_k = \sum_{i=1}^{k} X_i$, and $E_n = \max_{1 \leq k \leq n} S_k$. Using the Lévy inequality

$$\mathbb{P}\left(E_n \geq x\right) \leq 2\mathbb{P}(S_n \geq x) \Rightarrow$$
$$M(n) = \int_0^\infty \mathbb{P}\left(E_n \geq x\right) \mathrm{d}x \leq \int_0^\infty \mathrm{erfc}\left(\frac{x}{\sqrt{2n}}\right) \mathrm{d}x = \sqrt{\frac{2n}{\pi}}.$$

5 Discussion

The (DN) problem for monotone signals was first discussed in Donoho et al. (2013) in the context of monotone regression without regularization. There an upper bound was derived and the relation of the minimax error with the PTC for the (CS) problem was established. For the CS problem Pnevmatikakis and Paninski (2013) examined, in the context of sparse deconvolution, the case of signals where $\boldsymbol{x}(i+1) - \gamma \boldsymbol{x}(i)$ is sparse and non-negative, with $0 < \gamma < 1$ and close to 1, and identified the best, average, and worst cases depending on the location of the change points, without deriving a closed form expression. To the best of our knowledge, this paper presents for the first time an non-asymptotic closed form expression that captures the dependence on both the number and the location of the change points, and also characterizes the optimal regularizer. Future work includes the case of non-monotone sparsely varying signals, with the TV norm acting as the structure inducing function. The striking resemblance between the average SD (6) and the PTC for the case of non-negative sparse signals (Donoho et al., 2009), motivates a comparison between the average SD for this case and the PTC for recovering sparse signals using the l_1 norm. While a closed form solution for the SD is not available, some upper bounds appear in Cai and Xu (2015), simulations suggest a close match (Fig. 3).

6 Acknowledgements

Part of the work was performed while the author was with the Department of Statistics, Columbia University, NewYork, NY 10027. The author thanks L. Paninski, M. McCoy and J. Tropp for useful discussions.

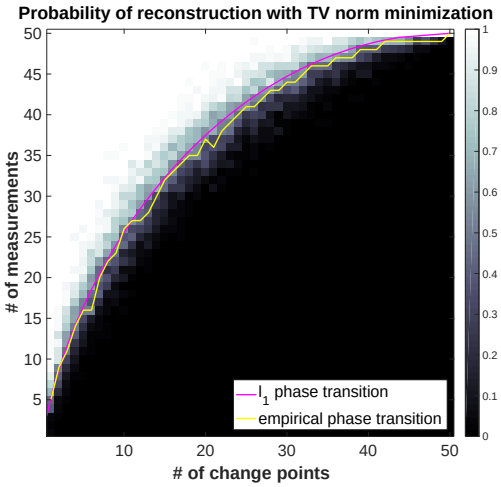
Probability of reconstruction with TV norm minimization
# of measurements
# of change points
l1 phase transition
empirical phase transition

Figure 3: Empirical calculation of reconstruction probability for sparsely varying signals. 50-dimensional piecewise constant signals were constructed with variable number of change points k and locations chosen uniformly at random. For each signal a random Gaussian sensing matrix was constructed with variable number of rows (measurements) m. Reconstruction was attempted by minimizing the TV norm subject to the measurements, and for each pair (k, m), 50 iterations were performed. The probability of success (color coded in the background) undergoes a phase transition. The empirical 50% success line (yellow) lies very close to the PTC for sparse signals (magenta) as is theoretically computed in Donoho et al. (2009).

References

Amelunxen, D., M. Lotz, M. B. McCoy, and J. A. Tropp (2014). Living on the edge: Phase transitions in convex programs with random data. Information and Inference, iau005.

Cai, J.-F. and W. Xu (2015). Guarantees of total variation minimization for signal recovery. Information and Inference, iav009.

Donoho, D., I. Johnstone, and A. Montanari (2013). Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising. IEEE Trans. Informat. Theory 59(6), 3396–3433.

Donoho, D., A. Maleki, and A. Montanari (2009). Message-passing algorithms for compressed sensing. Proceedings of the National Academy of Sciences 106(45), 18914.

Oymak, S. and B. Hassibi (2012). On a relation between the minimax risk and the phase transitions of compressed recovery. In Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on, pp. 1018–1025. IEEE.

Pnevmatikakis, E. and L. Paninski (2013). Sparse nonnegative deconvolution for compressive calcium imaging: algorithms and phase transitions. In Advances in Neural Information Processing Systems, Volume 26, pp. 1250–1258.